

Dynamic Auctions for On-Demand Services

Enrique Campos-Náñez, Natalia Fabra, and Alfredo Garcia

Abstract—We consider a market in which a finite number of firms compete in prices for the incoming demand for service. Upon every customer arrival, an independent auctioneer gathers bids from each one of the competing queuing systems and assigns the incoming customer to the system that submitted the lowest bid. We provide a simple characterization of Markov Perfect equilibrium in terms of “indifference prices”, i.e., price levels at which players are indifferent between committing available capacity or withholding it. We identify sufficient conditions for socially efficient performance in equilibrium.

Index Terms—Queuing Systems, Auctions, Markov Perfect Equilibrium, Dynamic Games.

I. INTRODUCTION

The pervasive nature of information technology (IT) in modern economies has inevitably changed the underlying structure through which economic transactions take place. As electronic commerce emerges as a viable retail channel, many firms have begun experimenting with alternative trading mechanisms such as auctions, guaranteed purchase contracts, group purchasing, etc. These new trading mechanisms offer unprecedented opportunities for improving the operational efficiency of capacity-constrained industries. Typically, the on-line trader or auctioneer is a “retail consolidator”, i.e., a firm that either buys cheaply, in advance, major blocks of capacity (e.g. airline seats or hotel rooms), or has access to a surplus inventory of deeply discounted capacity to be traded at possibly

marked-up prices. Often, retail consolidators or aggregators sell available capacity “on-demand”, i.e., an auction (or other trading mechanism) is launched upon request by a potential customer. In more traditional retail channels, potential customers incur “search” costs in obtaining and processing information about available prices. Similarly, sellers face costs associated with advertising, labor etc., typical of traditional retail activities. By relying on “retail consolidators” and their chosen trading mechanisms, firms can adequately price-in installed capacity while being able to focus on their core business competency. Customers may also save on the information search costs since the more successful on-line trading sites are exactly those that develop a reputation for enabling better price discovery. To summarize, on-demand trading has the potential to lower transaction costs and to enable better tracking between market prices and available capacity.

Our interest in auction-based mechanisms for on-demand service is also motivated by recent proposals for the implementation of a new paradigm known as “utility” computing. This label is generally used to refer to various schemes in which the computing resources of possibly different firms or “application service providers” are pooled in order to be allocated and priced upon customer request. It is believed that the utility computing paradigm would induce a more efficient exploitation of the existing IT infrastructure. In addition, by outsourcing IT services on-demand, small users avoid the substantial investment costs associated with owning and operating a separate IT division within their organizations.

In this paper, we explore the foundations of a theoretical model that captures many relevant features of an auction-based structure for on-demand trading of homogeneous

E. Campos-Náñez is with the Engineering Management and Systems Engineering Department, The George Washington University. N. Fabra is with the Universidad Carlos III. A. García is with the Department of Systems and Information Engineering, University of Virginia.

Manuscript received March 20, 2005; revised October 12, 2005.

services. Specifically, we consider a market in which a finite number of firms compete in prices for attending incoming demand for service. Each firm owns a queueing system (server with finite buffer) to provide the service. These queueing systems operate in parallel, and upon every customer arrival, an independent auctioneer gathers bids from each one of the competing queueing systems and assigns the incoming customer to the system that submitted the lowest bid.

Since only prices (and not capacity) can be updated in the short-run, we focus our attention on strategic pricing. In particular, we characterize the opportunity costs associated to committing available capacity at a given point in time. These opportunity costs are *strategic* in that they depend not only upon a firm's available capacity at a given point in time, but also on competitors' available capacity and the way these competitors price-in scarcity. Opportunity costs are also *dynamic* in that they depend on the "state" of the available capacity in the system. Hence, this setting is not one of repeated auctions, but one of dynamic auctions since the underlying opportunity costs that the players are facing vary over time.

The structure of the paper is the following: in section 2, we briefly review the related literature. In sections 3 and 4, we introduce the basic setup and the notion of Markov Perfect equilibrium. In section 5, we give a simple characterization of equilibrium in terms of "indifference prices", i.e., price levels at which players are indifferent between committing available capacity or withholding it. In section 6, we study conditions under which system-wide performance in equilibrium is efficient.

II. LITERATURE REVIEW

Our analysis is related to two strands of the literature: the literature that analyzes strategic behavior in queueing systems, and the auction literature.

The literature on strategic behavior in queueing systems has been recently surveyed in Hassin and Haviv (2002). In the pioneering work of Naor (1969), arriving customers are

assumed to be able to choose whether to join or leave depending on their expected disutility of delay. A congestion externality is therefore characterized as the increased delay due to other users' utilization of the service facility. If a customer enters the service system, then other customers may have to wait longer. Naor's work introduced the use of congestion pricing methods to alleviate the harmful effects of the congestion externality. Notable extensions to Naor's work can be found in Lippman and Stidham (1977) and Mendelson and Whang (1990).

Our analysis departs from this line of work in that we are interested in modeling competition among servers. On Chapter 8 of their recent survey, Hassin and Haviv (2002) provide a survey of published literature on competition among servers. Some papers (see Kalai et al.(1990) and Li (1992)) model non-price competition, (e.g. competition in service rates) while a few others, such as Lederer and Li (1997) study how delay costs affect both prices and operating policies. In all of these papers, firms maximize average reward. Consequently, equilibrium analysis is firmly grounded on (long-run) average reward and costs functions.

Our analysis differs substantially from the above mentioned papers in that, in our model, firms do account for the time value of money when evaluating their decisions. Consequently, the short-run performance of their queueing systems defines total discounted profits. Thus, while models that assume firms maximize average reward may be better suited to evaluate the long-run effects of say, installed capacity and/or service rates, models with discounted profit capture the need for dynamically adjusted prices to account for varying opportunity costs.

Our work is also related to the auction literature (see Klemperer (1999) for a complete survey of the auction literature). Several studies presenting overviews of the business to customers (B2C) online auctions have appeared since the debut of online auctions in 1995. Lucking-Reiley (2000) analyzes 142 B2C auctions. He finds that the types of auction mechanisms used on the Internet

are actually limited and traditional: the English, Vickrey, Dutch, and first-price sealed-bid auctions are virtually the only mechanisms used, and the English auction is, by far, the most common. Most auction models start by asserting assumptions on the valuation structure of bidders. In our setting, players valuations are not independent, as they are endogenously determined as a result of market interaction. Endogenous valuations arise because players' profits are interrelated across periods through costs - i.e., since buffer sizes are limited, serving the current customer may imply an infinite cost of serving the next incoming customer - and through prices - the prices that the rivals will be willing to charge to the future incoming customers will depend on whether their servers are idle or busy and on how they anticipate future customers will be allocated among players. However, we assume that given a pricing strategy, every player is able to determine the opportunity costs of his/her opponents. In this sense, our model is one with *complete* information but *endogenous* valuations.

III. SETUP

We now present the basic setup for price-based competition amongst queueing systems.

A.1 We assume customers arrive according to a Poisson process with rate λ . The service tasks requested by incoming customers are homogeneous and their willingness to pay, which we shall denote by v , is deterministic. There are n parallel queues (server plus buffer) owned by n different players. We shall denote by μ_i and K_i , the mean service rate (assuming service times are *i.i.d.* and exponentially distributed) and buffer capacity of player i 's queue. The above summarized structure is assumed to be common knowledge.

A.2 Auctions take place every time customers arrive. Given that arrivals constitute a Poisson process, the probability that two or more consumers arrive at the same time is zero. Hence, when an auction takes place, players will be bidding to service one customer

only. Bids can take any value in $[0, v]$ and players with full buffer are not allowed to bid in¹. Given the bids by players $b=(b_1, b_2, \dots, b_n) \in B = \prod_i [0, v]$, we shall denote by $b_{[k]}$ the k th-lowest bid, where $[k]$ represents the index of the player submitting that bid. The winner of the auction will be the one that submitted the lowest bid (in case of a tie among the players with the lowest bids, the index $[1]$ is randomly assigned to one of them). The price for the service $p^*(b)$ will be set at the second lowest bid. Formally, player's i demand $D_i(b)$ and the service price are defined as follows,

$$D_i(b) = \begin{cases} 1 & \text{if } b_i = b_{[1]} \\ 0 & \text{otherwise,} \end{cases}$$

where ties are randomly broken in case several players bid $b_{[1]}$. Also,

$$p^*(b) = \begin{cases} b_{[2]} & \text{if } [2] \neq \emptyset \\ b_{[1]} & \text{otherwise.} \end{cases}$$

The assumption that customers are not sensitive to service quality will be revisited in section 7.

A.3 The "state" of the system is the number of customers at each buffer. The state space is therefore $\mathcal{X} = \prod_{i=1}^n \{0, 1, 2, \dots, K_i\}$. Players have complete information about the state of the system.

A. Stationary Markovian Pricing Strategies.

We shall restrict our attention to Markovian pricing strategies, i.e., strategies where bids are a function of the current state of the system. Moreover, we are interested in stationary (i.e., time-invariant) strategies.

For each player i , a pure Markovian pricing strategy is denoted by the mapping $\pi_i : \mathcal{X} \mapsto [0, v]$. A Markovian strategy combination, $\pi = (\pi_1, \pi_2, \dots, \pi_n)$, is a vector of Markovian pricing strategies for each player. We let $\Pi \subseteq (\mathcal{X} \times [0, v])^n$ represent the set of all Markovian (pure) strategy combinations.

¹This assumption could be rationalized as follows: if large penalties for defaulting on the auction outcome are implemented, players with full buffer will abstain from bidding.

Let us denote by $r_i^\pi(x)$, player i 's expected payoff in the auction played when the state of the system is x and players follow strategy combination π , i.e.,

$$r_i^\pi(x) = p^*(\pi(x))D_i(\pi(x)).$$

We limit our interest to the times at which a new customer arrives. Given that a customer just arrived, and conditional upon the next customer arrival occurring after t units of time, we shall denote by $\mathbf{Q}^{\pi,t}$ the one-step transition probability matrix. Specifically, given states $x = (x_1, x_2, \dots, x_n)$, and $y = (y_1, y_2, \dots, y_n) \in \mathcal{X}$, the probability of transitioning from x to y given that the next arrival will occur at time t , and that the prices prescribed by the strategy π are used, is given by

$$\mathbf{Q}^{\pi,t}(x, y) = \prod_{i=1}^n Q_i^{\pi,t}(x, y_i),$$

where $Q_i^{\pi,t} \in \mathbb{R}^{|\mathcal{X}| \times K_i}$ is a matrix in which the (x, y_i) -th entry defines the probability of transitioning from state x to any state whose i -th component is y_i , i.e.,

$$Q_i^{\pi,t}(x, y_i) = \begin{cases} \bar{Q}_i^t(x_i, y_i) & \text{if } D_i(\pi(x)) = 1 \\ \underline{Q}_i^t(x_i, y_i) & \text{otherwise} \end{cases}$$

and \bar{Q}_i^t is the transition probability matrix for player i 's state variable, if at present time he/she is to attend the current customer conditioned on the next customer arriving after t units of time, i.e.,

$$\bar{Q}_i^t(x_i, y_i) = \begin{cases} 0 & \text{if } y_i > x_i + 1 \\ e^{-\mu_i t} \frac{(\mu_i t)^{(x_i - y_i + 1)}}{(x_i - y_i + 1)!} & \text{if } 0 < y_i \leq x_i + 1 \\ \sum_{l > x_i} e^{-\mu_i t} \frac{(\mu_i t)^l}{l!} & \text{if } y_i = 0, \end{cases}$$

and, \underline{Q}_i^t is the one-step transition probability vector for player i 's state variable, if at present time he/she is *not* to attend the current customer, and the next customer arrives after t time units, i.e.,

$$\underline{Q}_i^t(x_i, y_i) = \begin{cases} 0 & \text{if } y_i \geq x_i + 1 \\ e^{-\mu_i t} \frac{(\mu_i t)^{(x_i - y_i)}}{(x_i - y_i)!} & \text{if } 0 < y_i \leq x_i \\ \sum_{l \geq x_i} e^{-\mu_i t} \frac{(\mu_i t)^l}{l!} & \text{if } y_i = 0. \end{cases}$$

With this notation in hand, the value function is denoted by the mapping $V^{\pi,t} : \mathcal{X} \mapsto \mathcal{R}^n$ and can be recursively

defined as follows:

$$V_i^{\pi,t} = r_i^\pi + e^{-\rho t} \mathbf{Q}^{\pi,t} V_i^\pi,$$

where r_i^π is the immediate reward, the term $e^{-\rho t} \mathbf{Q}^{\pi,t} V_i^\pi$ is the expected (discounted) continuation reward and $V_i^\pi = E[V^{\pi,t}]$.

1) *Example 1: Computing the transition probability matrix:* Suppose that there are $n = 2$ parallel queues (server plus buffer) owned by the two different players. Each server has exponential service rate μ , and buffer capacity of 1. Thus, for $i = 1, 2$,

$$\bar{Q}_i^t = \begin{pmatrix} 1 - e^{-\mu t} & e^{-\mu t} & 0 \\ \sum_{k>1} e^{-\mu t} \frac{(\mu t)^k}{k!} & e^{-\mu t} (\mu t) & e^{-\mu t} \\ \sum_{k>1} e^{-\mu t} \frac{(\mu t)^k}{k!} & e^{-\mu t} (\mu t) & e^{-\mu t} \end{pmatrix},$$

and

$$\underline{Q}_i^t = \begin{pmatrix} 1 & 0 & 0 \\ \sum_{k>0} e^{-\mu t} \frac{(\mu t)^k}{k!} & e^{-\mu t} & 0 \\ \sum_{k>1} e^{-\mu t} \frac{(\mu t)^k}{k!} & e^{-\mu t} (\mu t) & e^{-\mu t} \end{pmatrix}.$$

The (x_i, y_i) -th entry of matrix \bar{Q}_i^t (respectively \underline{Q}_i^t) denotes the probability of transitioning from state x_i to state y_i for player i if the current customer is assigned (not assigned) to player i .

B. Valuation

Let $M_i \subset \mathcal{X}$ be defined as $M_i = \{x \in \mathcal{X} \mid x_i < K_i\}$. In the following discussion we restrict our attention to stationary Markovian pricing strategies π_i in the reduced domain M_i , i.e., $\pi_i : M_i \subset \mathcal{X} \mapsto [0, v]$. This is done with no loss of generality, since at every stage game, players with a full buffer capacity (i.e., $x_i = K_i$) will not participate in the next auction.

If players are assumed to follow a Markovian pricing policy π , from an individual perspective, each player is faced at each state with the choice between selling one unit of capacity at the given price, or withholding one unit of capacity. In order to compute the optimal response,

we solve the dynamic programming recursive equations conditional upon the next customer arrival occurring in t units of time. For $x \in M_i$ and $b_i \in [0, v]$ we have

$$\hat{V}_i^{\pi,t}(x; b) = p^*(b_i, \pi_{-i}(x)) D_i(b_i, \pi_{-i}(x)) + e^{-\rho t} \sum_{x' \in \mathcal{X}} \mathbf{Q}^{\pi,t}(x, x') \hat{V}_i^{\pi}(x'), \quad (1)$$

$$\hat{V}_i^{\pi}(x) = \sup_{0 \leq b_i \leq v} E[\hat{V}_i^{\pi,t}(x; b_i)], \quad (2)$$

where $(b_i, \pi_{-i}(x))$ stands for the strategy combination that equals π except at state x where player i bids b_i .

Equation (1) determines the value of selling today at the given prices (by bidding b_i). Equation (2) summarizes the value of today's best decision.

IV. MARKOV PERFECT EQUILIBRIUM

We are interested in Markovian strategy combinations that have the following property: at every time period, for any given state, no player can do strictly better by choosing a different price than the one prescribed by the strategy combination under consideration. This concept formalizes a notion of recursive rationality, i.e., play prescribed by the strategies from any state off the equilibrium path must also be in equilibrium (see Fudenberg and Tirole (1993)). As a refinement of Nash equilibrium, this solution concept filters out all “non-credible” Nash equilibria, i.e., those equilibrium strategies supported upon the basis of irrational play off the equilibrium path. In light of this, the MPE solution concept has more predictive power than Nash equilibrium. As argued in Maskin and Tirole (2001), a second advantage of MPE pertains to the simplicity of Markovian strategies which substantially reduces the number of parameters to be estimated in dynamic econometric models.

Formally, a strategy combination π^* is a Markov Perfect Equilibrium (MPE) if and only if, for every player i , and every state $x \in \mathcal{X}$,

$$V_i^{\pi^*}(x) \geq V_i^{(\pi_i, \pi_{-i}^*)}(x)$$

for all $\pi_i \neq \pi_i^*$, where (π_i, π_{-i}^*) is the strategy combination with player i bidding according to π_i (instead of π_i^*).

A. Indifference Prices

To characterize players' best replies for a given strategy combination π , we now introduce the notion of *indifference prices*. Given state x , and conditional upon the next customer arrival occurring in t units of time, we shall denote it by $\tilde{p}_i^{\pi,t}(x)$.

Definition 1: The indifference price map $\tilde{p}_i^{\pi,t}(x) : \Pi \times \mathcal{X} \mapsto [0, v]$ is such that

$$\tilde{p}_i^{\pi,t}(x) + e^{-\rho t} \sum_{x' \in \mathcal{X}} \bar{Q}_i^t(x_i, x'_i) \underline{\mathbf{Q}}_{-i}^t(x, x') \hat{V}_i^{\pi}(x') = e^{-\rho t} \sum_{x' \in \mathcal{X}} \underline{Q}_i^t(x_i, x'_i) \mathbf{Q}_{-i}^{\pi,t}(x, x') \hat{V}_i^{\pi}(x'), \quad (3)$$

where $\underline{\mathbf{Q}}_{-i}^t(x, x') = \prod_{j=1, j \neq i}^n \underline{Q}_j^t(x_j, x'_j)$, and $\mathbf{Q}_{-i}^{\pi,t}(x, x') = \prod_{j=1, j \neq i}^n \mathbf{Q}_j^{\pi,t}(x_j, x'_j)$.

The *indifference price* for player i is the price that equates the value obtained by selling one unit of buffer capacity today (left hand side in (3)) and withholding that unit for future revenue (right hand side in (3)). We remark that $\tilde{p}_i^{\pi,t}(x)$ is non-negative. By withholding capacity, a firm ensures that its buffer capacity next time a customer arrives will be greater than if it had served the current customer. Since withholding capacity by a given firm causes its rivals to use up their capacity, rivals' buffers will be less (next time a customer arrives) than if that firm were to serve the current customer. Since a firm's value rises as its available capacity increases and rivals' capacities decrease, it must hold that the future value associated with withholding is greater than (or equal to) the future value associated with selling. Also, $\tilde{p}_i^{\pi,t}(x) \leq v$ since the difference in per unit value between selling and withholding can not exceed v . Last, the expected indifference prices are denoted,

$$\tilde{p}_i^{\pi}(x) = E[\tilde{p}_i^{\pi,t}(x)]. \quad (4)$$

1) *Example 2: Computing the indifference prices:* Suppose that there are two parallel queues owned by two

different players. Each server has exponential service rate μ and no buffer capacity. Further assume that consumers' willingness to pay is deterministic and equal to v .

In states $(1,0)$ or $(0,1)$, when only one of the servers is available, the "available" player has monopoly power. Consequently, dynamic pricing strategies in equilibrium must set prices (for the "available" player) equal to v . The only non-trivial pricing decision is associated with state $(0,0)$, i.e., both servers are idle. Assuming that equilibrium pricing strategies are symmetric, we denote by $V(x_1, x_2)$ the value function for player 1, when the state of the system is (x_1, x_2) under the (symmetric) equilibrium pricing strategy.

Conditional upon the next customer arrival occurring in t units of time, the value function for player 1 when both servers are busy, $V^t(1,1)$, satisfies the following equation:

$$\begin{aligned} V^t(1,1) &= e^{-\rho t}[(1 - e^{-\mu t})^2 V(0,0) \\ &\quad + (1 - e^{-\mu t})e^{-\mu t}[V(0,1) + V(1,0)] \\ &\quad + e^{-2\mu t}V(1,1)] \end{aligned} \quad (5)$$

Given that both servers are busy, there are no immediate payoffs but only future expected payoffs that are discounted $e^{-\rho t}$. The expectation of future payoffs is obtained after considering the four possible transitions: service times for both busy servers are less than interval length t , thus the system state transitions to $(0,0)$; service time for one server is less than t , while the other server's service time exceeds t , hence, the transitions to asymmetric states $(0,1)$ and $(1,0)$; and last, service times for both busy servers exceed interval length t , thus the system state reverts to $(1,1)$.

It is easy to see that the value function for player 1 when its server is busy while that of player 2 is available, $V^t(1,0)$, is equal to $V^t(1,1)$, as there are no immediate payoffs to be made out of the current customer. Hence,

$$V^t(1,0) = V^t(1,1) \quad (6)$$

When player 1's server is available while player 2's is busy, player 1 earns a monopoly rent v on the current

customer as it faces no competition from its rival. Hence,

$$V^t(0,1) = v + V^t(1,1) \quad (7)$$

Last, when both servers are idle, player 1 would be indifferent between providing service at a price $\tilde{p}^t(0,0)$, or waiting (withholding capacity) if the following condition holds:

$$\begin{aligned} \tilde{p}^t(0,0) + e^{-\rho t}[e^{-\mu t}V(1,0) + (1 - e^{-\mu t})V(0,0)] \\ = e^{-\rho t}[e^{-\mu t}V(0,1) + (1 - e^{-\mu t})V(0,0)] \end{aligned}$$

Thus, $\tilde{p}^t(0,0) = e^{-(\rho+\mu)t}[V(0,1) - V(1,0)]$. From equations (6) and (7) we obtain, $\tilde{p}^t(0,0) = e^{-(\rho+\mu)t}v$. Its expected value is,

$$\begin{aligned} \tilde{p}(0,0) = E[\tilde{p}^t(0,0)] &= \int_0^\infty e^{-(\rho+\mu)t}v^* \lambda e^{-\lambda t} dt \\ &= \frac{\lambda}{\rho+\mu+\lambda}v \end{aligned}$$

Hence, the opportunity cost of one unit of buffer capacity, as measured by \tilde{p} , is increasing in λ and v , and decreasing in μ and ρ .

V. CHARACTERIZATION, EXISTENCE AND COMPUTATION OF MPE

Theorem 1: *A strategy combination of the form $\pi_i(x) = \tilde{p}_i^\pi(x)$ for all i and $x \in \mathcal{X}$, is a Markov Perfect Equilibrium.*

Proof: We first show that bidding $b_i < \tilde{p}_i^\pi(x)$ is weakly dominated by simply bidding $\tilde{p}_i^\pi(x)$. This deviation has two possible implications:

- a. If at state x , given bids (b_i, π_{-i}) , player i is the lowest bidder and the second lowest bid is below player i 's indifference price, i.e., $p^*(b_i, \pi_{-i}) < \tilde{p}_i^\pi(x)$, then player i 's expected payoff is

$$p^*(b_i, \pi_{-i}(x)) + E[e^{-\rho t} \sum_{x' \in \mathcal{X}} \underline{Q}_i^t(x_i, x'_i) \underline{Q}_{-i}^t(x, x') \hat{V}_i^\pi(x')] \quad (8)$$

On the other hand, if player i had bid at his expected indifference price, his expected payoff would have simply been,

$$E[e^{-\rho t} \sum_{x' \in \mathcal{X}} \underline{Q}_i^t(x_i, x'_i) \underline{Q}_{-i}^{\pi_{-i}}(x, x') \hat{V}_i^\pi(x')] \quad (9)$$

In the light of (3) and (4),

$$\begin{aligned}
& p^*(b_i, \pi_{-i}(x)) \\
& + E[e^{-\rho t} \sum_{x' \in \mathcal{X}} \bar{Q}_i^t(x_i, x'_i) \underline{Q}_{-i}^t(x, x') \hat{V}_i^\pi(x')] \\
& \leq p^*(\pi(x)) + E[e^{-\rho t} \sum_{x' \in \mathcal{X}} \bar{Q}_i^t(x_i, x'_i) \underline{Q}_{-i}^t(x, x') \hat{V}_i^\pi(x')] \\
& = E[e^{-\rho t} \sum_{x' \in \mathcal{X}} Q_i^t(x_i, x'_i) \mathbf{Q}_{-i}^{\pi, t}(x, x') \hat{V}_i^\pi(x')]
\end{aligned}$$

Hence, (8) does not exceed (9). In words, player i would have been better off bidding his indifference price.

- b. If by bidding $b_i < \tilde{p}_i^\pi(x)$, player i is the lowest bidder and the second lowest bidder is bidding above player i 's indifference price, player i would have been equally well-off bidding at his indifferent price.

Similarly, if by bidding $b_i < \tilde{p}_i^\pi(x)$ player i is *not* the lowest bidder, here again, player i would have been equally well-off bidding at his indifferent price

Second, we show that bidding $b_i > \tilde{p}_i^\pi(x)$ is weakly dominated by bidding his indifference price. If at state x the vector of bids (b_i, π_{-i}) is such that player i is *not* the lowest bidder, player i 's expected payoff is

$$E[e^{-\rho t} \sum_{x' \in \mathcal{X}} \underline{Q}_i^t(x_i, x'_i) \mathbf{Q}_{-i}^{\pi, t}(x, x') \hat{V}_i^\pi(x')] \quad (10)$$

whereas if player i had bid at his expected indifference price, his expected payoff would have been,

$$p^*(\pi(x)) + E[e^{-\rho t} \sum_{x' \in \mathcal{X}} \bar{Q}_i^t(x_i, x'_i) \underline{Q}_{-i}^t(x, x') \hat{V}_i^\pi(x')] \quad (11)$$

Again, in the light of (4) and the fact that $\tilde{p}_i^\pi(x) \leq p^*(\pi(x))$, we have that

$$\begin{aligned}
& E[e^{-\rho t} \sum_{x' \in \mathcal{X}} \underline{Q}_i^t(x_i, x'_i) \mathbf{Q}_{-i}^{\pi, t}(x, x') \hat{V}_i^\pi(x')] \\
& = \tilde{p}_i^\pi(x) + E[e^{-\rho t} \sum_{x' \in \mathcal{X}} \bar{Q}_i^t(x_i, x'_i) \underline{Q}_{-i}^t(x, x') \hat{V}_i^\pi(x')] \\
& \leq p^*(\pi(x)) + E[e^{-\rho t} \sum_{x' \in \mathcal{X}} \bar{Q}_i^t(x_i, x'_i) \underline{Q}_{-i}^t(x, x') \hat{V}_i^\pi(x')]
\end{aligned}$$

Hence, (10) does not exceed (11). In words, player i would have been better off bidding his indifference price. ■

In Theorem 1 we have shown that a strategy combination in which firms bid their expected indifference prices is a Markov Perfect Equilibrium. In other words, a fixed point of the expected indifference price operation is an equilibrium. A question remains on whether such a fixed point exists. In our next result we show that this is indeed the case.

Theorem 2: *Assuming $n=2$, there exists a strategy combination of the form $\pi_i(x) = \tilde{p}_i^\pi(x)$ for all i and $x \in \mathcal{X}$.*

Proof: We shall prove the expected indifference price operator is monotone in $\Pi^* \subset [0, v]^{|\mathcal{X}| \times n}$, where

$$\Pi^* = \{\pi \in [0, v]^{|\mathcal{X}| \times n} \mid \pi_i(x') \geq \pi_i(x),$$

for all i and $x', x \in \mathcal{X}$ such that $x' \geq x$.

Let us consider the indifference prices associated with a finite number of arrivals, say T . If $\pi \leq \pi'$ and $\pi, \pi' \in \Pi^*$, it follows that for fixed b_i ,

$$\begin{aligned}
& p^*((b_i, \pi_{-i}(x)) D_i(b_i, \pi_{-i}(x))) \\
& \leq p^*((b_i, \pi'_{-i}(x)) D_i(b_i, \pi'_{-i}(x)))
\end{aligned} \quad (12)$$

The value associated with the last arrival is

$$\hat{V}_i^\pi(x; T, T) = \sup_{0 \leq b_i \leq v} \{p^*((b_i, \pi_{-i}(x)) D_i(b_i, \pi_{-i}(x)))\}.$$

In light of (12),

$$\hat{V}_i^\pi(x; T, T) \leq \hat{V}_i^{\pi'}(x; T, T). \quad (13)$$

The value $\hat{V}_i^\pi(x; T-1, T)$ associated with the last two arrivals is given by the supremum over bids $b_i \in [0, v]$ of

$$\begin{aligned}
& p^*((b_i, \pi_{-i}(x)) D_i(b_i, \pi_{-i}(x))) \\
& + E[e^{-\rho t} \sum_{x' \in \mathcal{X}} \mathbf{Q}^{\pi, t}(x, x') \hat{V}_i^\pi(x'; T, T)].
\end{aligned} \quad (14)$$

In an analogous fashion to Theorem 1's proof, it can be seen that bidding $E[\tilde{p}_i^{\pi, t}(x; T-1, T)]$ where

$$\begin{aligned}
& \tilde{p}_i^{\pi, t}(x; T-1, T) \\
& + e^{-\rho t} \sum_{x' \in \mathcal{X}} \bar{Q}_i^t(x_i, x'_i) \underline{Q}_{-i}^t(x, x') \hat{V}_i^\pi(x'; T, T) \\
& = e^{-\rho t} \sum_{x' \in \mathcal{X}} Q_i^t(x_i, x'_i) \mathbf{Q}_{-i}^{\pi, t}(x, x') \hat{V}_i^\pi(x'; T, T),
\end{aligned}$$

is the optimal solution to (14). Given (13), we conclude

$$\begin{aligned} & \sum_{x' \in \mathcal{X}} [Q_i^t(x_i, x'_i) \mathbf{Q}_{-i}^{\pi, t}(x, x') - \bar{Q}_i^t(x_i, x'_i) \underline{\mathbf{Q}}_{-i}^t(x, x')] \\ & \quad \times \hat{V}_i^\pi(x'; T, T) \\ & \leq \sum_{x' \in \mathcal{X}} [Q_i^t(x_i, x'_i) \mathbf{Q}_{-i}^{\pi', t}(x, x') - \bar{Q}_i^t(x_i, x'_i) \underline{\mathbf{Q}}_{-i}^t(x, x')] \\ & \quad \times \hat{V}_i^{\pi'}(x'; T, T). \end{aligned}$$

Or equivalently,

$$\tilde{p}_i^{\pi, t}(x; T-1, T) \leq \tilde{p}_i^{\pi', t}(x; T-1, T).$$

Moreover, $\tilde{p}_i^{\pi, t}(x; T-1, T) \in \Pi^*$ and

$$\hat{V}_i^\pi(x; T-1, T) \leq \hat{V}_i^{\pi'}(x; T-1, T).$$

By finite induction we can show that

$$\tilde{p}_i^{\pi, t}(x; 0, T) \leq \tilde{p}_i^{\pi', t}(x; 0, T)$$

In the limit, $\tilde{p}_i^{\pi, t}(x) =$

$$\lim_{T \rightarrow \infty} \tilde{p}_i^{\pi, t}(x; 0, T) \leq \lim_{T \rightarrow \infty} \tilde{p}_i^{\pi', t}(x; 0, T) = \tilde{p}_i^{\pi', t}(x)$$

Finally, the result follows by invoking Tarski's fixed point theorem for the monotone map \tilde{p}^π on the complete lattice Π^* . ■

Theorem 2 motivates a very simple algorithm for the computation of MPE. Let $\pi_i^0(x) = v$ for all i and $x \in \mathcal{X}$. The algorithm's basic iteration is defined as follows:

$$\pi_i^{k+1}(x) = \tilde{p}_i^{\pi^k}(x)$$

By monotonicity, the sequence $\{\pi^k : k \geq 0\}$ is monotone decreasing with zero as lower bound. Therefore, there exists a limit point $\pi^* = \lim_{k \rightarrow \infty} \pi^k$. Finally, it follows that such limit point is a fixed point of the indifference price map, i.e.,

$$\pi_i^*(x) = \tilde{p}_i^{\pi^*}(x),$$

for all i and $x \in \mathcal{X}$.

A. First-Price Auction

As a Corollary to Theorem 1, we consider now the case when the auction format is slightly altered so that the auction winner is paid according to his/her bid. In practice,

the first-price auction rule is more widely used than the second-price format. In this setting, the player with the lowest expected indifference price has an incentive to “undercut” the player with the second lowest expected indifference price. We shall assume, as is standard in the literature, that in case of bidding ties, demand is served by the player with the lower expected indifference price (in case of ties among players with equal indifference prices, we shall assume that both players face an equal probability of serving demand). For a given strategy combination π and a given state x , the optimal “undercutting” price is given by,

$$\underline{p}^\pi(x) = \arg \max_{p \leq \tilde{p}_{[2]}^\pi(x) \leq v} (p - \tilde{p}_{[1]}^\pi(x))$$

Corollary 1: Given $x \in \mathcal{X}$ and π , let $[k]$ denote the index associated with the k -th lowest expected indifference price. If π is of the form,

$$\begin{aligned} \pi_{[1]}(x) &= \underline{p}^\pi(x), \\ \pi_i(x) &= \tilde{p}_i^\pi(x), \quad i \neq [1], \end{aligned}$$

then π is a Markov Perfect Equilibrium.

Proof: By construction, the player with lowest indifference price, i.e., $\tilde{p}_{[1]}^\pi(x)$, maximizes expected markup over its indifference price by bidding $\underline{p}^\pi(x)$. Since $\underline{p}^\pi(x) \leq \tilde{p}_{[2]}^\pi(x)$ all the other players have no incentive to deviate. ■

VI. EFFICIENCY

We are interested in studying whether the Markov Perfect Equilibrium characterized above induces a system-wide efficient utilization of resources. Since an incoming customer can only be rejected if all the buffers are busy, a system-wide efficient routing policy is one that maximizes the expected number of customers served over an unbounded horizon. Hordijk and Koole (1992) have shown that any system-wide efficient policy has the following structure: an arriving customer should be assigned to a faster server when that server has a shorter queue. They refer to this type of policy as the “Shorter Queue Faster Server Policy, SQFSP”. As it is pointed out in Hordijk and Koole (1992), this characterization is incomplete in

general, but sufficient whenever the servers have identical service time distributions. In this case, the optimal policy follows a simple ‘‘Shortest Queue’’ rule; incoming customers should be routed to the shortest queue. An application to Hordijk and Koole’s characterization is the following result:

Corollary 2: *Assume expected indifference prices are monotone in buffer state i.e.,*

$$x_i \geq x_j \implies \tilde{p}_i^\pi(x) \geq \tilde{p}_j^\pi(x),$$

where $\pi_i(x) = \tilde{p}_i^\pi(x)$ for all i and $x \in \mathcal{X}$. With identical service time distributions, π induces a system-wide efficient routing policy.

Nonetheless, when servers have different service time distributions, the performance of the auction-based controlling mechanism may be inefficient as we shall illustrate in our next example.

A. Example 3: the MPE may not be efficient

Again, suppose that there are two bufferless queues owned by two different players and that consumers’ willingness to pay is deterministic and equal to v . At state $(1, 1)$ requests will be denied, while at states $(0, 1)$ and $(1, 0)$ the requests will be accepted at a price v . Therefore, the only interesting state is $(0, 0)$. As in example 2 above, it is possible to state a set of equations for the total discounted value for each player. We shall denote by $V_i^t(x_1, x_2)$ player i ’s discounted value in equilibrium when the state of the system is (x_1, x_2) and the next customer arrival occurs in t units of time;

$$\begin{aligned} V_1^t(1, 0) &= e^{-\rho t}[(1 - e^{-\mu_1 t})(1 - e^{-\mu_2 t})V_1(0, 0) \\ &\quad + e^{-(\mu_1 + \mu_2)t}V_1(1, 1) + \\ &\quad + e^{-\mu_1 t}(1 - e^{-\mu_2 t})V_1(1, 0) \\ &\quad + (1 - e^{-\mu_1 t})e^{-\mu_2 t}V_1(0, 1)] \end{aligned}$$

$$\begin{aligned} V_1^t(0, 1) &= v + e^{-\rho t}[(1 - e^{-\mu_1 t})(1 - e^{-\mu_2 t})V_1(0, 0) \\ &\quad + e^{-(\mu_1 + \mu_2)t}V_1(1, 1) \\ &\quad + e^{-\mu_1 t}(1 - e^{-\mu_2 t})V_1(1, 0) \\ &\quad + (1 - e^{-\mu_1 t})e^{-\mu_2 t}V_1(0, 1)] \\ V_1^t(1, 1) &= e^{-\rho t}[(1 - e^{-\mu_1 t})(1 - e^{-\mu_2 t})V_1(0, 0) \\ &\quad + e^{-(\mu_1 + \mu_2)t}V_1(1, 1) \\ &\quad + e^{-\mu_1 t}(1 - e^{-\mu_2 t})V_1(1, 0) \\ &\quad + (1 - e^{-\mu_1 t})e^{-\mu_2 t}V_1(0, 1)] \end{aligned}$$

Finally, we write down the equation for the indifference price at state $(0, 0)$:

$$\begin{aligned} \tilde{p}_1^t(0, 0) + e^{-\rho t} [(1 - e^{-\mu_1 t})V_1(0, 0) + e^{-\mu_1 t}V_1(1, 0)] \\ = e^{-\rho t} [(1 - e^{-\mu_2 t})V_1(0, 0) + e^{-\mu_2 t}V_1(1, 0)] \end{aligned}$$

Conversely, for player two, we have the following equations:

$$\begin{aligned} V_2^t(1, 0) &= v + e^{-\rho t}[(1 - e^{-\mu_1 t})(1 - e^{-\mu_2 t})V_2(0, 0) \\ &\quad + e^{-(\mu_1 + \mu_2)t}V_2(1, 1) \\ &\quad + e^{-\mu_1 t}(1 - e^{-\mu_2 t})V_2(1, 0) \\ &\quad + (1 - e^{-\mu_1 t})e^{-\mu_2 t}V_2(0, 1)] \\ V_2^t(0, 1) &= e^{-\rho t}[(1 - e^{-\mu_1 t})(1 - e^{-\mu_2 t})V_2(0, 0) \\ &\quad + e^{-(\mu_1 + \mu_2)t}V_2(1, 1) \\ &\quad + e^{-\mu_1 t}(1 - e^{-\mu_2 t})V_2(1, 0) \\ &\quad + (1 - e^{-\mu_1 t})e^{-\mu_2 t}V_2(0, 1)] \\ V_2^t(1, 1) &= e^{-\rho t}[(1 - e^{-\mu_1 t})(1 - e^{-\mu_2 t})V_2(0, 0) \\ &\quad + e^{-(\mu_1 + \mu_2)t}V_2(1, 1) \\ &\quad + e^{-\mu_1 t}(1 - e^{-\mu_2 t})V_2(1, 0) \\ &\quad + (1 - e^{-\mu_1 t})e^{-\mu_2 t}V_2(0, 1)] \end{aligned}$$

together with the indifference price equation:

$$\begin{aligned} \tilde{p}_2^t(0, 0) + e^{-\rho t} [(1 - e^{-\mu_2 t})V_2(0, 0) + e^{-\mu_2 t}V_2(1, 0)] \\ = e^{-\rho t} [(1 - e^{-\mu_1 t})V_2(0, 0) + e^{-\mu_1 t}V_2(1, 0)] \end{aligned}$$

At state $(0, 0)$, the player with the lowest indifference price wins the auction at a price equal to the highest indifference price. If we posit that $\tilde{p}_1^t(0, 0) < \tilde{p}_2^t(0, 0)$ (i.e., the winner

at state $(0,0)$ is player 1 and the price for service is $\tilde{p}_2^t(0,0)$ two additional equations are added:

$$\begin{aligned} V_1^t(1,0) &= \tilde{p}_2^t(0,0) + e^{-\rho t} \\ &\quad \times [(1 - e^{-\mu_1 t})V_1(0,0) + e^{-\mu_1 t}V_1(1,0)] \\ V_2^t(0,0) &= e^{-\rho t} [(1 - e^{-\mu_1 t})V_2(0,0) + e^{-\mu_1 t}V_2(1,0)] \end{aligned}$$

Solving numerically the above equations for $\tilde{p}_1^t(0,0)$ and $\tilde{p}_2^t(0,0)$ with the parameter values $v = 1$, $\rho = 0.9$, $\lambda = 1$ and $\mu_1 = 0.75$, the conjecture $\tilde{p}_1^t(0,0) < \tilde{p}_2^t(0,0)$ is verified whenever $\mu_2 > 0.75 = \mu_1$. In conclusion, the player with the highest service rate will also have the highest indifference price due to the higher prospect of future monopoly rents. This results in inefficiency since by assigning arrivals in state $(0,0)$ to the server with the lowest service rate, there is higher probability of service rejections.

VII. INCORPORATING DELAY SENSITIVITY

Up to this point, our model assumes that customers are not sensitive to expected delays. While this could apply to settings in which the disutility of delay is negligible when compared to the surplus associated with a service completion, it is of interest to extend our setup to incorporate delay sensitive customers. In this section, we provide a numerical illustration of the effects associated to incorporating time sensitive customers.

A. Setup Revisited

We now revisit assumption A.2 in the basic setup for price-based competition amongst queueing systems as presented in section 3.

A.2' Auctions take place every time customers arrive.

Arriving customers experience a delay cost c per unit of time. Therefore, given a state $x \in \mathcal{X}$, and bids $b = (b_1, b_2, \dots, b_n) \in B = \Pi_i[0, v]$, we define [1] to be the index associated with the best offer, i.e.,

$$[1] = \min_i \left\{ b_i + \frac{cx_i}{\mu_i} \right\}$$

where ties are randomly broken. Player's i demand $D_i(b, x)$ is defined as follows,

$$D_i(b, x) = \begin{cases} 1 & \text{if } i = [1] \\ 0 & \text{otherwise,} \end{cases}$$

The price for service $p^*(b)$ is defined as follows

$$p^*(b) = \begin{cases} b_{[2]} & \text{if } [2] \neq \emptyset \\ b_{[1]} & \text{otherwise.} \end{cases}$$

where

$$[2] = \min_{i \neq [1]} \left\{ b_i + \frac{cx_i}{\mu_i} \right\}.$$

Given the Markovian structure of pricing policies, the definition of indifferent prices given in equation (3) remains unaltered. The interplay between different pricing strategies and the expected delays is captured by the value function definition in (1) and (2).

B. Numerical Illustration

Consider a realization of the queueing system, defined by the sequence of states $\{x^1, x^2, \dots\}$, with $x^k \in \mathcal{X}$ for all $k = 1, 2, \dots$ occurring at times $\{t_1 < t_2 < \dots\}$ obtained through simulation of policy $\pi \in [0, v]^{|\mathcal{X}| \times n}$ with initial state x^0 . Let $\{k_m\}_{m=1}^\infty$ be the subsequence of system states corresponding to a new arrival. Given a roll-out horizon $\tau > 0$, we define the rollout estimate of the value function for state $x = x^0 \in \mathcal{X}$ as

$$\tilde{V}_i^\pi(x, \tau) = \sum_{j=1}^{\tau} e^{-\rho(t_{k_j})} \gamma_i^\pi(x^{k_j}).$$

Using these approximations, player i can estimate its indifference price by letting

$$\hat{p}_i^\pi(x) = \tilde{V}_i^\pi(x, \tau) - \tilde{V}_i^\pi(x + e_i, \tau),$$

where e_i is the unit vector in the i -th dimension, defined for all states $x_i < K_i$.

Let $\pi^0 \in [0, v]^{|\mathcal{X}| \times n}$ be an initial bidding policy, and define recursively

$$\pi^{k+1}(x) = \max\{0, \min\{v, (1 - \gamma_k)\pi^k(x) + \gamma_k \hat{p}_i^{\pi^k}(x)\}\}, \quad (15)$$

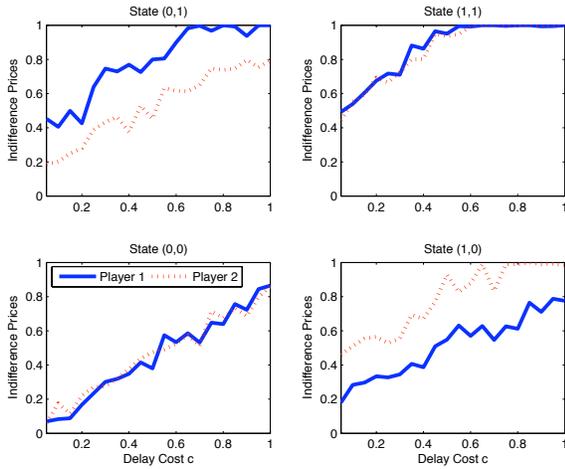


Fig. 1. The plots illustrate the evolution of the indifference prices for all states of the system for the case $K_1 = K_2 = 2$, for values of the delay cost ranging from $c = [.05, .1, \dots, 1]$. States where any of the two players are full have been omitted, since for those states the indifference prices are v .

for $k = 1, 2, \dots$, and the standard conditions $\sum_k \gamma_k = \infty$, and $\sum_k \gamma_k^2 < \infty$ (e.g. $\gamma_k = \frac{1}{k}$). Clearly, any fixed point of this recursion corresponds to an MPE.

We applied our algorithm to a symmetrical system with $K_1 = K_2 = 2$, an arrival rate of $\lambda = 1$, and service rates of $\mu_1 = \mu_2 = 3/4$, and a roll-out horizon $\tau = 100$. In Fig. 1, we present the a summary of different runs for values of the delay cost c ranging from $[.05, .10, \dots, .95, 1.0]$. For each value of c , one hundred iterations of the roll-out algorithm were carried out, and the last set of indifferent prices, i.e., π^{100} , is shown in the plots. The experiment reflects the symmetry of the setup, and also shows that although price increases with the delay cost, the MPE is not necessarily efficient. In a similar experiment on an asymmetrical system, with $K_1 = 2$, and $K_2 = 4$, we obtain the results shown in Figs. 3-??, which show similar tendencies. In Fig. 2 we show the results of 100 iterations of the algorithm.

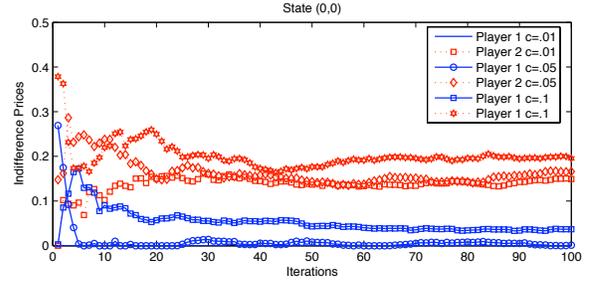


Fig. 2. Evolution the estimated indifference prices for state $(0, 0)$ with $K_1 = 2, K_2 = 4$, and $c \in \{.01, .05, .1\}$.

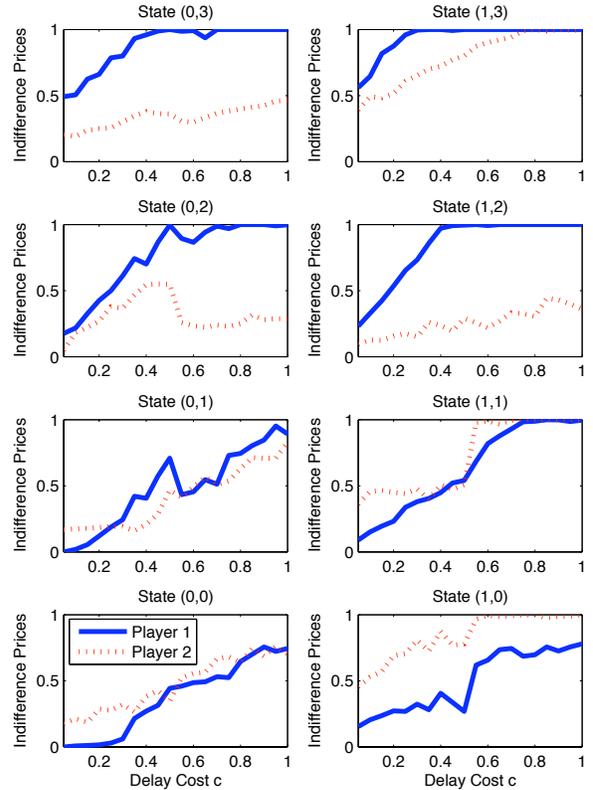


Fig. 3. The plots illustrate the evolution of the indifference prices for all states of the system for the case $K_1 = 2, K_2 = 4$, for delay costs ranging in $c = [.05, .1, \dots, 1.0]$.

VIII. CONCLUSIONS

In this paper, we have developed a theoretical model that captures many relevant features of an auction-based structure for on-line trading of homogeneous services. A simple characterization of Markov Perfect equilibrium in terms of “indifference prices”, i.e., price levels at which players are indifferent between committing available capacity or withholding it from the market, reflects the assessment of opportunity costs associated to available capacity at a given point in time. Interestingly, strategic effects are non-negligible since these opportunity costs depend not only on each firm’s available capacity at a given point in time, but also on competitors’ available capacity. Therefore, we have a market in which a firm pricing strategy determines partially its competitors costs. This feature inevitably affects social efficiency. We give a sufficient condition under which system-wide performance in equilibrium is efficient. The incorporation of delay-sensitive demand and relaxing the assumption on complete information about the state of the system are topics of further research.

REFERENCES

- [1] Fudenberg, D., Tirole, J. Game Theory. MIT press, 1993.
- [2] Gagliano R., Fraser M. and Schaefer S. “Auction allocation of computing resources”, Communications of the ACM, Vol. 38 (6) pp. 88 - 102, 1995.
- [3] Hassin R., Haviv, M. “To Queue or Not to Queue: Equilibrium Behavior in Queueing Systems”. Kluwer Academic Publishers, 2002.
- [4] Hordijk, A., Koole, G. “On the Assignment of Customers to Parallel Queues”. Probability in the Engineering and Informational Sciences, Vol. 6, pp. 495-511, 1992.
- [5] Kalai, E., Kamienn, M., Rubinovitch, M. “Optimal Service Speeds in a Competitive Environment”. Management Science, Vol. 38 pp. 1147-1190, 1990.
- [6] Klemperer, P. “Auction Theory: A Guide to the Literature”. Journal of Economic Surveys, Vol. 13 (3), pp. 227-286, 1999.
- [7] Li, L. “The Role of Inventory in Delivery Price Competition”. Management Science, Vol. 38, pp. 192-204, 1992.
- [8] Lederer, P., Li, L.. “Pricing, Production, Scheduling and Delivery-Time Competition”, Operations Research, Vol. 45, pp. 407-420, 1997.
- [9] Lippman, S., Stidham, S. “Individual vs Social Optimization in Exponential Congestion Systems”. Operations Research, Vol. 25 pp. 233-247, 1977.
- [10] Lucking-Reiley, D. “Auctions on the Internet: Whats being auctioned, and how?” Journal of Industrial Economics. 48(3) 227-252, 2000.
- [11] Maskin E. and Tirole J., “Markov Perfect Equilibrium I. Observable Actions”. Journal of Economic Theory, Vol. 100 (2) pp.191-219, 1997.
- [12] Masuda, Y., Whang S. “Dynamic Pricing for Network Service: Equilibrium and Stability”. Management Science, Vol. 45, pp. 867-869, 1999.
- [13] Mendelson, H. H., Whang, S. “Optimal Incentive Compatible Priority Pricing for the $M/M/1$ Queue”. Operations Research, Vol. 38, pp. 870-883, 1990.
- [14] Naor, P. “The Regulation of Queue Size by Levying Tolls. Econometrica”, Vol.37, No. 1, pp. 15-24, 1969.
- [15] Vickrey, W. “Counterspeculation, Auctions, and Competitive Sealed Tenders”, Journal of Finance, Vol. 16, pp. 8-37, 1961.
- [16] Yen, B.P.-C. Wu, O.Q. “Internet scheduling environment with market-driven agents”, IEEE Transactions on Systems, Man and Cybernetics, Part A, Vol. 34 (2) pp. 281 - 289, 2004.